

Intelligent Virtual Patients for Training Clinical Skills

Albert “Skip” Rizzo, Patrick Kenny, and Thomas D. Parsons*

*Institute for Creative Technologies
University of Southern California, Los Angeles, CA, USA
email: arizzo@usc.edu
email: {kenny, tparsons} @ict.usc.edu

Abstract

The article presents the design process of intelligent virtual human patients that are used for the enhancement of clinical skills. The description covers the development from conceptualization and character creation to technical components and the application in clinical research and training. The aim is to create believable social interactions with virtual agents that help the clinician to develop skills in symptom and ability assessment, diagnosis, interview techniques and interpersonal communication. The virtual patient fulfills the requirements of a standardized patient producing consistent, reliable and valid interactions in portraying symptoms and behaviour related to a specific clinical condition.

Keywords: Simulation, Virtual Humans, Virtual Training, Virtual Patients, Behaviour Generation, Speech Synthesis, Dialogue Generation, speech recognition

Digital Peer Publishing Licence

Any party may pass on this Work by electronic means and make it available for download under the terms and conditions of the current version of the Digital Peer Publishing Licence (DPPL). The text of the licence may be accessed and retrieved via Internet at <http://www.dipp.nrw.de/>.

First presented at the Virtual Reality International Conference VRIC 2009, extended and revised for JVRB

1 Introduction

Over the last 15 years, a virtual revolution has taken place in the use of simulation technology for clinical purposes. Technical advances in the areas of computation speed and power, graphics and image rendering, display systems, tracking, interface technology, haptic devices, authoring software and artificial intelligence have supported the creation of low-cost and usable PC-based Virtual Reality (VR) systems. At the same time, a determined and expanding cadre of researchers and clinicians, have not only recognized the potential impact of VR technology, but have now generated a significant research literature that documents the many clinical targets where VR can add value over traditional assessment and intervention approaches [Hol05, PR08, PE08, RBR05, Riv05]. This convergence of the exponential advances in underlying VR enabling technologies with a growing body of clinical research and experience has fueled the evolution of the discipline of Clinical Virtual Reality. And this state of affairs now stands to transform the vision of future clinical practice and research in the disciplines of psychology, medicine, neuroscience, physical and occupational therapy, and in the many allied health fields that address the rehabilitative needs of those with clinical disorders.

A short list of areas where Clinical VR has been usefully applied includes fear reduction with phobic clients [PR08, PE08], treatment for Post Traumatic Stress Disorder [RRG⁺09], stress management in cancer patients [SPPA⁺04], acute pain reduction during wound care and physical therapy with burn patients [HPM⁺04], body image disturbances in patients with eating disorders [Riv05], navigation and spatial training in children and adults with motor impairments [SFW98, RSKM04], functional skill train-

ing and motor rehabilitation with patients having central nervous system dysfunction (e.g., stroke, TBI, SCI cerebral palsy, multiple sclerosis, etc.) [Hol05, WKFK06] and in the assessment (and in some cases, rehabilitation) of attention, memory, spatial skills and executive cognitive functions in both clinical and unimpaired populations [RBR05, RK05, RSKM04]. To do this, VR scientists have constructed virtual airplanes, skyscrapers, spiders, battlefields, social settings, beaches, fantasy worlds and the mundane (but highly relevant) functional environments of the schoolroom, office, home, street and supermarket. These efforts are no small feat in light of the technological challenges, scientific climate shifts and funding hurdles that many researchers have faced during the early development of this emerging technology.

Concurrent with the emerging acknowledgement of the unique value of Clinical VR by scientists and clinicians, has come a growing awareness of its potential relevance and impact by the general public. While much of this recognition may be due to the high visibility of digital 3D games, the Nintendo Wii, and massive shared internet-based virtual worlds (World of Warcraft, Halo and 2nd Life), the public consciousness is also routinely exposed to popular media reports on clinical and research VR applications. Whether this should be viewed as “hype” or “help” to a field that has had a storied history of alternating periods of public enchantment and disregard, still remains to be seen. Regardless, growing public awareness coupled with the solid scientific results have brought the field of Clinical VR past the point where skeptics can be taken seriously when they characterize VR as a “fad technology”.

These shifts in the social and scientific landscape have now set the stage for the next major movement in Clinical VR. With advances in the enabling technologies allowing for the design of ever more believable context-relevant “structural” VR environments (e.g. homes, classrooms, offices, markets, etc.), the next important challenge will involve *populating* these environments with virtual human (VH) representations that are capable of fostering believable interaction with *real* VR users. This is not to say that representations of human forms have not usefully appeared in Clinical VR scenarios. In fact, since the mid-1990’s, VR applications have routinely employed VHs to serve as stimulus elements to enhance the realism of a virtual world simply by their static presence.

For example, VR exposure therapy applications

have targeted simple phobias such as fear of public speaking and social phobia using virtual social settings inhabited by “still-life” graphics-based characters or 2D photographic sprites [AZHR05, PSB02, Kli05]. By simply adjusting the number and location of these VH representations, the intensity of these anxiety-provoking VR contexts could be systematically manipulated with the aim to gradually habituate phobic patients and improve their functioning in the real world. Other clinical applications have also used animated graphic VHs as stimulus entities to support and train social and safety skills in persons with high functioning autism [RCN⁺03, PSC06] and as distracter stimuli for attention assessments conducted in a virtual classroom [PBBR07, RKM⁺06].

In an effort to further increase the pictorial realism of such VHs, Virtually Better Inc., began incorporating whole video clips of crowds into graphic VR fear of public speaking scenarios [Vir10]. They later advanced the technique by using blue screen captured video sprites of humans inserted into graphics-based VR social settings for social phobia and cue exposure substance abuse treatment applications. The sprites were drawn from a large library of blue-screen captured videos of actors behaving or speaking with varying degrees of provocation. These video sprites could then be strategically inserted into the scenario with the aim to modulate the emotional state of the patient by fostering encounters with these 2D video VH representations.

The continued quest for even more realistic simulated human interaction contexts led other researchers to the use of panoramic video [MPWR07, RGG03] capture of a real world office space inhabited by hostile co-workers and supervisors to produce VR scenarios for anger management research. With this approach, the VR scenarios were created using a 360-degree panoramic camera that was placed in the position of a worker at a desk and then actors walked into the workspace, addressed the camera (as if it was the targeted user at work) and proceeded to verbally threaten and abuse the camera, vis-à-vis, the worker. Within such photorealistic scenarios, VH video stimuli could deliver intense emotional expressions and challenges with the aim of the research being to determine if this method could engage anger management patients to role-play a more appropriate set of coping responses.

However, working with such fixed video content to foster this form of *faux* interaction or exposure has sig-

nificant limitations. For example, it requires the capture of a large catalog of possible verbal and behavioral clips that can be tactically presented to the user to meet the requirements of a given therapeutic approach. As well, this fixed content cannot be readily updated in a dynamic fashion to meet the challenge of creating a real time interaction with a patient with the exception of only very constrained social interactions. Essentially, this process can only work for clinical applications where the only requirement is for the VH character to deliver an open-ended statement or question that the user can react to, but is lacking in any truly fluid and believable interchange following a response by the user. This has led some researchers to consider the use of artificially intelligent VH agents as entities for simulating human-to-human interaction in virtual worlds.

Clinical interest in artificially intelligent agents designed for interaction with humans can trace its roots to the work of MIT AI researcher, Joe Weizenbaum. In 1966, he wrote a language analysis program called ELIZA that was designed to imitate a Rogerian therapist. The system allowed a computer user to interact with a virtual therapist by typing simple sentence responses to the computerized therapist's questions. Weizenbaum reasoned that simulating a non-directional psychotherapist was one of the easiest ways of simulating human verbal interactions and it was a compelling simulation that worked well on teletype computers (and is even instantiated on the internet today; http://www-ai.ijs.si/eliza-cgi-bin/eliza_script). In spite of the fact that the illusion of Eliza's intelligence soon disappears due to its inability to handle complexity or nuance, Weizenbaum was reportedly shocked upon learning how seriously people took the ELIZA program [HM00]. And this led him to conclude that it would be immoral to substitute a computer for human functions that "...involves interpersonal respect, understanding, and love." [Wei76].

More recently, seminal research and development has appeared in the creation of highly interactive, artificially intelligent (AI) and natural language capable virtual human agents. No longer at the level of a prop to add context or minimal faux interaction in a virtual world, these VH agents are designed to perceive and act in a 3D virtual world, engage in face-to-face spoken dialogues with real users (and other VHs) and in some cases, they are capable of exhibiting human-like emotional reactions. Previous classic work on virtual humans in the computer graphics community

focused on perception and action in 3D worlds, but largely ignored dialogue and emotions. This has now changed. Both in appearance and behavior, VHs have now "grown up" and are ready for service in a variety of clinical and research applications.

Artificially intelligent VH agents can now be created that control computer generated bodies and can interact with users through speech and gesture in virtual environments [GJA⁺02]. Advanced virtual humans can engage in rich conversations [TMG⁺08], recognize nonverbal cues [MdKG08], reason about social and emotional factors [GM04] and synthesize human communication and nonverbal expressions [TMMK08]. Such fully embodied conversational characters have been around since the early 90's [BC05] and there has been much work on full systems to be used for training [Eva89, KHG⁺07, PI04, RGH⁺01], intelligent kiosks [MD06], and virtual receptionists [BSBH06].

The focus of this paper will be on our work developing and evaluating the use of AI VHs for use as virtual patients for training clinical skills in novice clinicians. While we believe that the use of VHs to serve the role of virtual therapists is still fraught with both technical and ethical concerns [RSR02], we have had success in the initial creation of VHs that can mimic the content and interaction of a patient with a clinical disorder for training purposes and we will describe our technical approach and clinical application research.

2 Intelligent Virtual Human Patients

An integral part of medical and psychological clinical education involves training in symptom or ability assessment, diagnosis, interviewing skills and interpersonal communication. Traditionally, medical students and other healthcare practitioners initially learn these skills through a mixture of classroom lectures, observation and role-playing practice with standardized patients—persons recruited and trained to take on the characteristics of a real patient, thereby affording students a realistic opportunity to practice and be evaluated in a simulated clinical environment. Although a valuable tool, there are limitations with the use of standardized patients that can be mitigated through VR simulation technology. First, standardized patients are expensive. For example, although there are 130 medical schools in the U.S., only five sites provide standardized patient assessments as part of the U.S. Medical Licensing Examination at a cost of several thou-

sand dollars per student [Edu10]. Second is the issue of standardization. Despite the expense of standardized patient programs, the standardized patients themselves are typically low skilled actors making about \$10/hr and administrators face constant turnover resulting in considerable challenges to the consistency of patient portrayals. This limits the value of this approach for producing reliable and valid interactions needed for the psychometric evaluation of clinicians in training. Finally, the diversity of the conditions that standardized patients can characterize is limited by availability of human actors and their skills. This is even a greater problem when the actor needs to be a child, adolescent, elder, or in the mimicking of nuanced or complex symptom presentations.

In this regard, Virtual Patients (VPs) can fulfill the role of standardized patients by simulating a particular clinical presentation with a high degree of consistency and realism [SHJ⁺06], as well as being always available for anytime-anywhere training. There is a growing field of researchers applying VP's to training and assessment of bioethics, basic patient communication, interactive conversations, history taking, clinical assessments, and clinical decision making [BG06, BPPO07, LFR⁺06, KRP⁺07, PKN⁺08]. Initial results suggest that VPs can provide valid and reliable representations of live patients [TFK⁺06, RJD⁺07]. Similar to the compelling case made over the years for Clinical VR generally [RK05], VP applications can likewise enable the precise stimulus presentation and control (dynamic behavior, conversational dialog and interaction) needed for rigorous laboratory research, yet embedded within the context of ecologically relevant environments [KRP⁺07, PKN⁺08, RJD⁺07].

The virtual patients described in this paper focus on characters primarily with psychological disorders. There has already been a significant amount of work creating virtual patients and virtual dummies that represent physical problems [KLD⁺08], such as how to attend to a wound, or interviewing a patient with a stomach problem to assess the condition [LFR⁺06]. However, research into the use of VPs in psychology and related psychosocial clinical training has been limited [BG06, KRP⁺07, BSW⁺06, CBB⁺98, DJR⁺05, JDR⁺06]. Beutler and Harwood [BH04] describe the development of a VR system for training in psychotherapy and summarize training-relevant research findings. We could not find reference to any other use of VPs in a psychotherapy course to date, despite online searches through MEDLINE, Ovid, and the psy-

chotherapy literature. From this, it is our view that the design of intelligent VPs that have realistic and consistent human-to-human interaction and communication skills would open up possibilities for clinical psychosocial applications that address interviewing skills, diagnostic assessment and therapy training.

3 Designing a Virtual Patient Domain

The VP domains for the clinical applications of interest for our research derives from a cognitive behavioral framework that focuses on a dual assessment of the patient's description of objective behaviors/symptoms and the VPs cognitive interpretations that co-exist in the appraisal of such behaviors/symptoms. This mental health domain offers some interesting opportunities and challenges for both the design of the characters as well as in the design of a training system that aims to enhance interviewing, differential diagnosis, and therapeutic communication skills. The domain also offers a framework for a wide variety of modeling issues including: verbal and non-verbal behavior for the characters, cognition, affect, rational and irrational behavior, personality, and for symptoms seen with psychopathology. Designing virtual patients that act as humans is a process that requires a thorough understanding of the domain, the technology and how to apply the technology to the domain. The process to create virtual characters is broken down into five steps as described here and illustrated in Figure 2.

3.1 Conceptualization

The first step is to define the concept and goals of the project, in this case being the application of virtual human technology for training optimal interaction with patients exhibiting specific clinical conditions. A close collaboration with experts in fields related to the specific clinical domain is required to define the typical symptoms, goals and interactional features of the VP. The concept is further refined via role-playing exercises, clinical manuals, and by soliciting knowledge from subject matter experts on what constitutes a useful interview with a various patient types. The methodology used in our approach is an iterative process of data collecting, testing, and subsequent refinement. While the preliminary goal of the project is to use VPs for individual diagnostic interview training,

the eventual goal is to also have VPs utilized in classroom and other group settings.

3.2 Design

After the concept has been established, the next step involves the process of designing the age, gender and appearance of the VP character, their background and clinical history, the kinds of actions and dialog they need to produce, and the artwork for the environment where the character will reside. The current virtual environment for our work was modeled after a typical clinician's office and was designed to represent a "safe" setting for the conduct of a clinical interview.

The artwork plays a crucial role in defining the characters behavior, attitude and condition. One of our initial projects involved the development of a VP named "Justine", who was the victim of a sexual assault (see Figure 1). However, we wanted to create an adolescent female, which while relevant for this clinical situation, was general enough such that the artwork would not be exclusively tied to any one clinical condition. For this particular VP we wanted a typical teenager with a T-shirt and blue jeans. It was a careful process to design this character, since subtle characteristics such as a belly shirt or ripped jeans, might provoke questions that would need to be accounted for in the dialog domain. A minimalist approach for initial steps in this design process was envisioned to support bootstrapping development of later character design as the system was evolved based on user input.

3.3 Development

The development of a VP domain is an iterative process that can be constrained by the technology. For example, in our system there was a specific way to encode the dialog in the system and rules that govern the behavior of the character. This began with the encoding of a corpus of questions that would likely be asked by a clinician and responses that could be given by the VP using a language model in the speech recognizer. The current system uses pre-recorded speech for the character and all of the many VP responses require recording by a voice actor. Following repeated tests with pilot users, the language model is updated with question and answer variations added into the system. Again, this is a bootstrapping process that in most cases requires many iterative cycles to create a model of natural and believable therapeutic interaction.



Figure 1: Adolescent Female Virtual Patient ("Justine")

3.4 Application

Once the VP system is constructed and internally tested with pilot participants, it is then brought to the clinical test site, where a larger sample of clinical professionals can be employed for experimental tests and application development purposes. In current tests of the virtual patient application, various general technical questions have been addressed relating to the accuracy of the speech recognizer, VP response relevance and general system performance. Non-technical evaluation involves evaluation of the breadth and specificity of assessment questions asked, investigation of the balance of clinician questions and statements (e.g. information gathering, interpretations, reflections, etc.), use of follow-up questions to refine initial diagnostic impressions, and general assessment of user ratings and feedback regarding usability and believability of the scenario.

3.5 Evaluation and Training Interview Skills

The art and science of evaluating interviewing skills using VPs is still a young discipline with many challenges. One formative approach is to compare performances obtained during interviews with both live standardized patients and with VPs, and then to conduct correlational analyses of metrics of interest. This information can then be evaluated relative to an Objective Structured Clinical Examination (OSCE) [HSDW75, WOR05]. Such tests typically take from 20-30 minutes and require a faculty member to watch the student perform a clinical interview while being videotaped. The evaluation consists of a

self-assessment rating along with faculty assessment and a review of the videotape. This practice is common, although is applied variably, based on the actors, available faculty members and space and time constraints at the training site. A general complication involved in teaching general interviewing skills is that there are multiple theoretical orientations and techniques to choose from and the challenge will be to determine what commonality exists across these methods for the creation of usable and believable VPs that are adaptable to all clinical orientations. To minimize this problem in our initial efforts, we are concentrating on assessing the skills required to diagnose very specific mental disorders (i.e., conduct disorder, PTSD, depression, etc.). We also use the setting of an initial intake interview to constrain the test setting for acquiring comprehensible data. While the clinician may have some knowledge of why the patient is there (i.e., a referral question), they need to ask the patient strategic questions to obtain a detailed history useful for specifying a clinical condition in support of coming to a differential diagnosis and for formulating a treatment plan. In this manner, the system is designed to allow novice clinicians the opportunity to practice asking interview questions that eventually lead to the narrowing down of the alternative diagnostic options, leading to the arrival of a working diagnosis based on the VP meeting the criteria for a specific Diagnostic and Statistical Manual of Mental Disorders IV-TR (DSM) diagnosis (or not).

4 Virtual Human Architecture

The VP system is based on our existing virtual human architecture [KHG⁺07, SGJ⁺06]. The general architecture supports a wide range of virtual humans from simple question/answering to more complex ones that contain cognitive and emotional models with goal-oriented behavior. The architecture is a modular distributed system with many components that communicate by message passing. Because the architecture is modular it is easy to add, replace or combine components as needed. For example in the virtual human architecture, the natural language section is divided into three components: a part to understand the language, a part to manage the dialog and a part to generate the output text. This is all combined into one statistical language component for the current VP system.

Interaction with the system works as follows and can be seen in Figure 3. A user talks into a microphone

that records the audio signal that is sent to a speech recognition engine. The speech engine converts that into text. The text is then sent to a statistical response selection module. The module picks an appropriate verbal response based on the input text question. The response is then sent to a non-verbal behavior generator that selects animations to play for the text, based on a set of rules. The output is then sent to a procedural animation system along with a pre-recorded or a generated voice file. The animation system plays and synchronizes the gestures, speech and lip-syncing for the final output to the screen. The user then listens to the response and asks more questions to the character.

4.1 Human Speech Input

A human user talks to the system using a head-mounted close-capture USB microphone. The user's speech is converted into text by an automatic speech recognition system. We use the SONIC speech recognition engine from the University of Colorado, Boulder [Pel01] customized with acoustic and language models for the domain of interest [SGN05]. In general a language model is tuned to the domain word lexicon that helps improve performance; however the speech recognition system is speaker independent and does not need to be trained beforehand. User's voice data is collected during each testing session, to enable the collection of words not recognized by the speech recognizer and further enhance the lexicon. The speech recognition engine processes the audio data and produces a message with the surface text of the user's utterance that is sent to the question / response selection module.

4.2 Question / Response Selection

To facilitate rapid development and deployment of virtual human agents, a question/response module was developed for natural language processing called the NPCEditor [LPT⁺06], a user friendly software module for editing questions and responses, binding responses to questions and training and running different statistical text classifiers. The question/response selection module receives a surface text message from the speech recognition module, analyzes the text, and selects the most appropriate response. This response selection process is based on a statistical text classification approach developed by the natural language group at the USC Institute for Creative Technologies [LM06]. The approach requires a domain de-

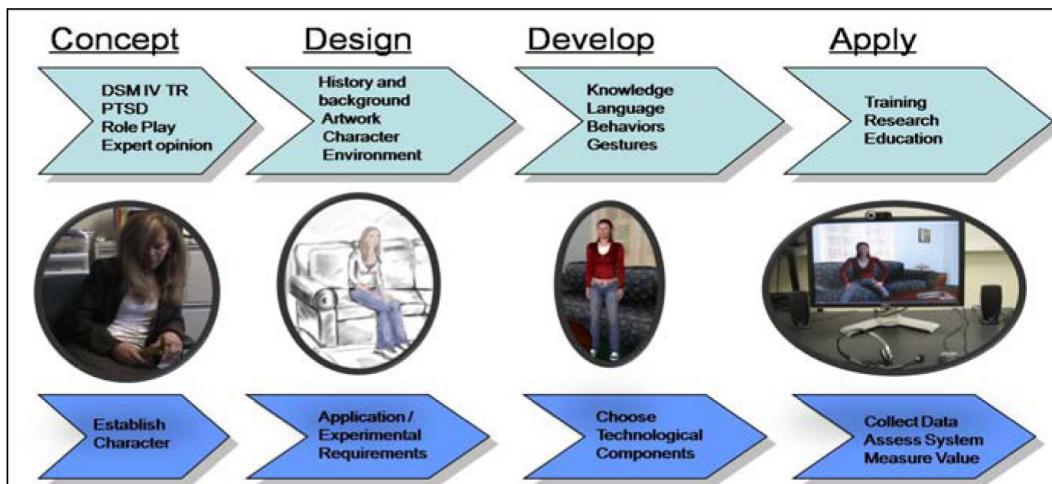


Figure 2: Development process for creating a Virtual Patient

signer to provide some sample questions for each response. There is no limit to the number of answers or questions, but it is advised to have at least four to ten questions for each answer. When a user question is sent from the speech recognition module, the system uses the mapping between the answers and sample questions as a “dictionary” to translate the question into a representation of a “perfect” answer. It then compares that representation to all known text answers and selects the best match. This approach was developed for the Sgt. Blackwell and Sgt. Star virtual human projects [KHG⁺07] and has been shown to outperform traditional state-of-the-art text classification techniques. Responses can be divided into several categories based on the type of action desired. To make the character appear more realistic and interactive it is advisable to have responses in each of the categories. The category types are as follows:

- *On-topic* - These are answers or responses that are relevant to the domain of the conversation. These are the answers the system has to produce when asked a relevant question. Each on-topic answer should have a few sample questions and single sample questions can be linked to several answers. The text classifier generally returns a ranked list of answers and the system makes the final selection based on the rank of the answer and whether the answer has been used recently. That way if the user repeats his questions, he may get a different response from the system.
- *Off-topic* - These are answers for questions that do not have domain-relevant answers. They can

be direct, e.g., “I do not know the answer”, or evasive, e.g., “I will not tell you” or “Better ask somebody else”. When the system cannot find a good on-topic answer for a question, it selects one of the off-topic lines.

- *Repeat* - If the classifier selects an answer tagged with this category, the system does not return that answer but replays the most recent response. Sample questions may include lines like “What was that?” or “Can you say that again?” Normally, there is at most one answer of this category in the domain answer set.
- *Alternative* - If the classifier selects an answer tagged with this category, the system attempts to find an alternative answer to the most recent question. It takes the ranked list of answers for the last question and selects the next available answer. Sample questions may include lines like “Do you have anything to add?” Normally, there is at most one answer tagged with this category in the answer set.
- *Pre-repeat* - Sometimes the system has to repeat an answer. For example, this happens when a user repeats a question and there is only one good response available. The system returns the same answer again but indicates that it is repeating itself by playing a pre-repeat-tagged line before the answer, e.g., “I told you already.” There is no need to assign sample questions to these answer lines.

- *Delayed* - These are the lines from the system that prompt the user to ask about a domain related thing, e.g., “Why dont you ask me about” Such a response is triggered if the user asks too many off-topic questions. The system would return an off-topic answer followed by a delayed-tagged answer. That way the system attempts to bring the conversation back into the known domain. This category has no sample questions assigned.

Once the output response is selected, it is packaged up into a Functional Markup Language (FML) message structure. FML allows the addition of elements such as affect, emphasis, turn management, or coping strategies. For the VP, the response selection module does not add any additional information besides the text.

4.3 Behavior Generation

The FML message is sent to the non-verbal behavior generator (NVBG) [LM06] that applies a set of rules to select gestures, postures and gazes for the virtual character. Since the VPs in these applications are sitting down in a clinical office, the animations mainly consist of arm movements, wave offs and head shakes or nods. The VP character does not do any posture shifts or go into standing posture, although in the next version more gestures and postures will be added. Once the NVB selects the appropriate behavior for the input text, it then packages this up into a Behavioral Markup Language (BML) [KKM⁺06] structure and sends it to a procedural animation system for character action control.

4.4 Character Control and Output

This last part of the process is the execution and display of the characters’ multimodal behavior. This is accomplished with a procedural animation system called *Smartbody* [TMMK08]. *Smartbody* takes as input the BML message from the NVBG which contains the set of behaviors that need to be executed for the head, facial expressions, gaze, body movements, arm gestures, speech and lip syncing and synchronizes all of this together. These behaviors need to be in sync with the output speech to look realistic. *Smartbody* is capable of using generated or pre-recorded speech. Our current VP uses pre-recorded speech. *Smartbody* is hooked up to a visualization engine, in this case the

Unreal Tournament game engine for the graphics output. *Smartbody* controls the character in the game engine and also specifies which sound to play for the character speech output. *Smartbody* is also capable of controllers that perform specific actions based on rules, functions or timing information, such as head nods and gaze following. The controllers are seamlessly blended in with the input animations specified in the BML. The VP does not make extensive use of controllers, however, future work is to design controllers for certain behavior patterns such as gaze aversion or physical tics the character might have. A motex, which is a looping animation file, can be played for the character to give it a bit of ambient sway or motion, or in the VP case, finger tapping for more realistic idle behavior.

5 Summary of Virtual Patient Projects

In addition to providing a historical context for this research area, our presentation at VRIC will focus primarily on the results from our initial research with VPs modeled to emulate conduct disorder and PTSD diagnostic categories, briefly summarized below. We will also detail some new projects that involve the use of VPs to train assessment of parenting skills, suicide potential with a depressed patient, racial bias testing and a virtual elderly person with Alzheimers disease designed for caregiver training.

5.1 Adolescent Male with Conduct Disorder

Our initial project in this area involved the creation of a virtual patient, named “Justin”. The clinical attributes of Justin were developed to emulate a conduct disorder profile as found in the DSM IV-TR. Justin portrays a 16-year old male with a conduct disorder who is being forced to participate in therapy by his family (see Figure 4). Justin history is significant for a chronic pattern of antisocial behavior in which the basic rights of others and age-appropriate societal norms are violated. He has stolen, been truant, broken into someone’s car, been cruel to animals, and initiated physical fights. The system was designed to allow novice clinicians to practice asking interview questions, to create a positive therapeutic alliance and gather clinical information from this very challenging VP.

This VP was designed as a first step in our research. At the time, the project was unfunded and thus

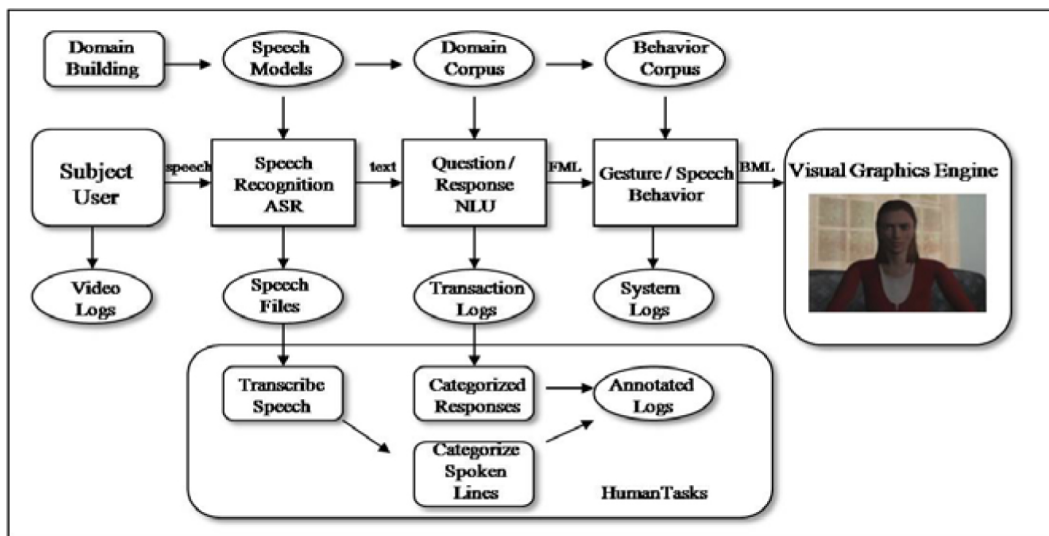


Figure 3: Virtual Patient Data Flow Diagram



Figure 4: Adolescent male Virtual Patient (Justin)

required us to take the economically inspired route of recycling a character from a military negotiation-training scenario to play the part of Justin. The research group agreed that this sort of patient was one that could be convincingly created within the limits of the technology (and funding) available to us at the time. For example, such resistant patients typically respond slowly to therapist questions and often use a limited and highly stereotyped vocabulary. This allowed us to create a believable VP within limited resources for dialog development. As well, novice clinicians have been typically observed to have a difficult time learning the value of “waiting out” periods of silence and non-participation with these patients.

We initially collected user interaction and dialog data from a small sample of psychiatric residents and

psychology graduate students as part of our iterative design process to evolve this application area. Assessment of the system was completed by 1) experimenter observation of the participants as they communicated with the virtual patient; and 2) questionnaires.

To adequately evaluate the system, we determined a number of areas that needed to be addressed by the questionnaires: 1) The behavior of the virtual patient should match the behavior one would expect from a patient in such a condition (e.g. verbalization, gesture, posture, and appearance); 2) Adequacy of the communicative discourse between the virtual patient and the participants; 3) Proficiency (e.g. clarity, pace, utility) of virtual patient discourse with the participant; and 4) Quality of the speech recognition of utterances spoken.

Although there were a number of positive results, the overall analysis revealed some shortcomings. First, positive results from the questionnaires were found to be related to overall system performance. Further, participants reported that the system 1) simulated the real-life experience (i.e. ranked 5 or 6 out of 7); and 2) the verbal and non-verbal behavior also ranked high (i.e. between 5 and 7). However, results also revealed that some participants found aspects of the experience “frustrating”. For example, some participants complained that they were receiving anticipated responses and the system tended to repeat some responses too frequently. This was due to the speech recognition’s inability to evaluate certain of the stimulus words. Further, there were too many “brush off” responses from the virtual patient when participant questions were

outside the virtual patient's dialog set. There was also a concern that participants ascribe characteristics to the virtual patient that in fact are not present. For example, although the virtual patient responded "Yes" to a question about whether the virtual patient "hurt animals", in actuality the system did not recognize the input speech. This may lead to confusion if the virtual patient responds inconsistently. In fact one of the most substantial lessons learned was the amount of conversation tracking needed for the topical questions asked by the participant to allow the virtual patient's responses to be consistent throughout the session.

The project produced a successful proof of concept demonstrator, which then led to the funding that currently supports our research in this area.

5.2 Adolescent Female with PTSD due Sexual Assault

Our 2nd VP project involved the creation of a female sexual assault victim, "Justina". The aim of this work was two fold: 1. Explore the potential for creating a system for use as a clinical interview trainer for promoting sensitive and effective clinical interviewing skills with a VP that had experienced significant personal trauma; and 2. By manipulating the dialog content in multiple versions of "Justina", provide a test of whether novice clinicians would ask the appropriate questions needed to determine if Justina reports symptoms that meet the criteria for the DSM-4 diagnosis of PTSD.

For the PTSD content domain, 459 questions were created that mapped roughly 4 to 1 to a set of 116 responses. The aim was to build an initial language domain corpus generated from subject matter experts and then capture novel questions from a pilot group of users (psychiatry residents) during interviews with Justina. The novel questions generated could then be fed into the system in order to iteratively build the language corpus. We also focused on how well subjects asked questions that covered the six major symptom clusters that can characterize PTSD following a traumatic event. While this did not give the Justina character a lot of depth, it did provide more breadth, which for initial testing seemed prudent for generating a wide variety of questions for the next Justina iteration.

In the initial test, a total of 15 Psychiatry residents (6 females, 9 males; mean age = 29.80, SD 3.67) participated in the study and were asked to perform a 15-minute interaction with the VP to take an initial his-

tory and determine a preliminary diagnosis based on this brief interaction with the character. The participants were asked to talk normally, as they would to a standardized patient, but were informed that the system was a research prototype that uses an experimental speech recognition system that would sometimes not understand them. They were free to ask any kind of question and the system would try to respond appropriately.

From post questionnaire ratings on a 7-point likert scale, the average subject rating for believability of the system was 4.5. Subjects reported their ability to understand the patient at an average of 5.1, but rated the system at 5.3 as frustrating to talk to, due to speech recognition problems, out of domain questions or inappropriate responses. However most of the participants left favorable comments that they thought this technology will be useful, they enjoyed the experience of trying different ways to talk to the character and attempting to get an emotional response for a difficult question. When the patient responded back appropriately to a question they found that very satisfying. Analysis of concordance between user questions and VP response pairs indicated moderate effects sizes for Trauma inquiries ($r = 0.45$), Re-experiencing symptoms ($r = 0.55$), Avoidance ($r = 0.35$), and in the non-PTSD general communication category ($r = 0.56$), but only small effects were found for Arousal/Hypervigilance ($r = 0.13$) and Life impact ($r = 0.13$). These relationships between questions asked by a novice clinician and concordant replies from the VP suggest that a fluid interaction was sometimes present in terms of rapport, discussion of the traumatic event, the experience of intrusive recollections and discussion related to the issue of avoidance.

This work has now formed the basis for a new project that will further modify the Justin and Justina VP characters for military clinical training. In one project, both Justin and Justina will appear as soldiers who are depressed and possibly contemplating suicide, as a training tool for teaching clinicians and other military personnel how to recognize the potential for this tragic event to occur. Another project focuses on the military version of "Justina" with the aim to develop a training tool that clinicians can practice therapeutic skills for addressing the growing problem of sexual assault within military ranks (See Figure 5). The system can also be used by command staff to foster better skills for recognizing the signs of sexual assault in subordinates under their command and for improving the



Figure 5: Military version of Justina



Figure 6: Justina experiencing VR Exposure Therapy

provision of support and care.

And finally, we intend to use both military VP versions in the role of a patient who is undergoing Virtual Reality exposure therapy (See Figure 6) for PTSD. The VP will appear in a simulation of a VR therapy room wearing a HMD, while the therapist practices the skills that are required for appropriately fostering emotional engagement with the trauma narrative as is needed for optimal therapeutic habituation.

This simulation of a patient experiencing VR exposure therapy uses the *Virtual Iraq/Afghanistan* PTSD system [RRG⁺09] as the VR *context*, and the training methodology is based on the Therapist's Manual created for that VR application by Rothbaum, Difede and Rizzo [RDR08]. We believe the "simulation of an activity that occurs within a simulation" is a novel concept that has not been reported previously in the VR literature!

If this exploratory work continues to show promise,

we intend to address a longer-term vision—that of creating a comprehensive DSM diagnostic trainer that has a diverse library of VPs modeled after each diagnostic category. The VPs would be created to represent a wide range of age, gender and ethnic backgrounds and could be interchangeably loaded with the language and emotional model of any of the DSM disorders. We believe this research vision will afford many research opportunities for investigating the functional and ethical issues involved in the process of creating and interacting with virtual humans. While ethical challenges may be more intuitively appreciated in cases where the target user is a patient with a clinical condition, the training of clinicians with VPs will also require a full appreciation of how this form of training impacts clinical performance with *real* patients.

References

- [AZHR05] P. L. Anderson, E. Zimand, L. F. Hodges, and B. O. Rothbaum, *Cognitive behavioral therapy for public-speaking anxiety using virtual reality for exposure*, *Depression and Anxiety* **22** (2005), no. 3, 156–158, ISSN 1091-4269.
- [BC05] T. Bickmore and J. Cassell, *Advances in natural multimodal dialogue systems*, ch. 2: Social Dialogue with Embodied Conversational Agents, Dordrecht: Springer, 2005.
- [BG06] Timothy Bickmore and Toni Giorgino, *Health Dialog Systems for Patients and Consumers*, *Journal of Biomedical Informatics* **39** (2006), no. 5, 556–571, ISSN 1532-0464.
- [BH04] Larry E. Beutler and Mark T. Harwood, *Virtual reality in psychotherapy training*, *Journal of clinical psychology* **60** (2004), no. 3, 317–330, ISSN 0021-9762.
- [BPPO07] Timothy W. Bickmore, Laura M. Pfeifer, and Michael K. Paasche-Orlow, *Health Document Explanation by Virtual Agents*, IVA '07 Proceedings of the 7th international conference on Intelligent Virtual Agents, Springer-Verlag, 2007, pp. 183–196, ISBN 978-3-540-74996-7.

- [BSBH06] Sabarish Babu, Stephen Schmugge, Tiffany Barnes, and Larry F. Hodges, *What Would You Like to Talk About? An Evaluation of Social Conversations with a Virtual Receptionist*, IVA Intelligent Virtual Agents (J. Gratch, Michael Young, Ruth Aylett, Daniel Ballin, and Patrick Olivier, eds.), Springer, 2006, Volume 4133, pp. 169–180, ISBN 978-3-540-37593-7.
- [BSW⁺06] Thomas Bernard, Amy Stevens, Peggy Wagner, Nicole Bernard, Christopher Oxendine, Kyle Johnsen and Robert Dickerson, Andrew Raji, Benjamin Lok, Margaret Duerson, Marc Cohen, Lori Schumacher, Garrett J. Harper, and D. Scott Lind, *A Multi-Institutional Pilot Study to Evaluate the Use of Virtual Patients to Teach Health Professions Students History-Taking and Communication Skills*, Simulation in Healthcare: The Journal of the Society for Simulation in Healthcare **1** (2006), no. 2, 92–93, ISSN 1559-2332, Abstracts Presented at the 6th Annual International Meeting on Medical Simulation.
- [CBB⁺98] J. Cassell, T. Bickmore, M. Billingham, L. Campbell, K. Chang, H. Vilhjálmsón, and H. Yan, *An Architecture for Embodied Conversational Characters*, Proceedings of the First Workshop on Embodied Conversational Characters, 1998.
- [DJR⁺05] R. Dickerson, K. Johnsen, A. Raji, B. Lok, J. Hernandez, and A. Stevens, *Evaluating a script-based approach for simulating patient-doctor interaction*, Proceedings of the International Conference of Human-Computer Interface Advances for Modeling and Simulation, 2005.
- [Edu10] Educational Commission for Foreign Medical Graduates, *ECFMG CSEC Centers*, www.ecfm.org/usmle/step2cs/centers.html, 2010, Last visited February 14th, 2011.
- [Eva89] David R. Evans, *Essential Interviewing: A Programmed Approach to Effective Communication*, 3rd ed., Brooks/Cole, 1989, ISBN 0-534-09960-2.
- [GJA⁺02] J. Gratch, Rickel J., E. Andre, Cassell J, E. Petajan, and N. Badler, *Creating Interactive Virtual Humans: Some Assembly Required*, IEEE Intelligent Systems **17** (2002), no. 4, 54–63, ISSN 1541-1672.
- [GM04] Jonathan Gratch and Stacy Marsella, *A domain independent framework for modeling emotion*, Cognitive Systems Research **5** (2004), no. 4, 269–306, ISSN 1389-0417.
- [HM00] Steve R. Howell and Rob Muller, *Computers in Psychotherapy: A New Prescription*, www.psychology.mcmaster.ca/beckerlab/showell/ComputerTherapy.PDF, 2000, Last visited February 11th, 2011.
- [Hol05] M. K. Holden, *Virtual Environments for Motor Rehabilitation: Review*, CyberPsychology and Behavior **8** (2005), no. 3, 187–211, ISSN 1094-9313.
- [HPM⁺04] Hunter G. Hoffman, David R. Patterson, Jeff Magula, Gretchen J. Carrougher, Karen Zeltzer, Stephen Dagadakis, and Sam R. Sharar, *Water-friendly virtual reality pain control during wound care*, J. Clin. Psychol. **60** (2004), no. 2, 189–195, ISSN 1097-4679.
- [HSDW75] R. M. Harden, M. Stevenson, W. W. Downie, and G. M. Wilson, *Assessment of clinical competence using objective structured examination*, British Medical Journal **1** (1975), no. 5955, 447451, ISSN 0267-0623.
- [JDR⁺06] Kyle Johnsen, Robert Dickerson, Andrew Raji, Cyrus Harrison, Benjamin Lok, Amy Stevens, and D. Scott Lind, *Evolving an immersive medical communication skills trainer*, Presence: Teleoperators and Virtual Environments **15** (2006), no. 1, 33–46, ISSN 1054-7460.
- [KHG⁺07] Patrick Kenny, Arno Hartholt, Jonathan Gratch, William Swartout, David Traum, Stacy Marsella, and Diane Piepol,

- Building Interactive Virtual Humans for Training Environments*, The Interservice/Industry Training, Simulation & Education Conference (IITSEC), 2007.
- [KKM⁺06] Stefan Kopp, Brigitte Krenn, Stacy Marsella, Andrew N. Marshall, Catherine Pelachaud, Hannes Pirker, Kristinn R. Thorisson, and Hannes Vilhjalmsson, *Towards a Common Framework for Multimodal Generation: The Behavior Markup Language*, IVA Intelligent virtual agents, Lecture Notes in Computer Science Vol. 4133, 2006, pp. 205–217, ISBN 3-540-37593-7.
- [KLD⁺08] A. Kotranza, B. Lok, Adeline Deladisma, Adeline Deladisma, and Carla M. Pugh, *Virtual Human + Tangible Interface = Mixed Reality Human. An Initial Exploration with a Virtual Breast Exam Patient*, Proceedings of IEEE VR 2008 Virtual Reality Conference, 2008, pp. 99–106, ISBN 978-1-4244-1971-5.
- [Kli05] E. Klinger, *Virtual Reality Therapy for Social Phobia: its Efficacy through a Control Study*, Paper presented at: Cybertherapy 2005. Basel, Switzerland, 2005.
- [KRP⁺07] Patrick Kenny, Albert A. Rizzo, Thomas D. Parsons, Jonathan Gratch, and William Swartout, *A Virtual Human Agent for Training Clinical Interviewing Skills to Novice Therapists*, Annual Review of Cybertherapy and Telemedicine **5** (2007), 82–90, ISSN 1554-8716.
- [LFR⁺06] Benjamin Lok, Richard E. Ferdig, Andrew Raij, Kyle Johnsen, Robert Dickerson, Jade Coutts, Amy Stevens, and D. Scott Lind, *Applying Virtual Reality in Medical Communication Education: Current Findings and Potential Teaching and Learning Benefits of Immersive Virtual Patients*, Virtual Reality **10** (2006), no. 3-4, 185–195, ISSN 1359-4338.
- [LM06] Jina Lee and Stacy Marsella, *Nonverbal Behavior Generator for Embodied Conversational Agents*, IVA Intelligent virtual agents, Lecture Notes in Computer Science Vol. 4133, 2006, pp. 243–255, ISBN 3-540-37593-7.
- [LPT⁺06] Anton Leuski, Jarrell Pair, David Traum, Peter J. McNERney, Panayiotis Georgiou, and Ronakkumar Patel, *How to talk to a hologram*, IUI '06 Proceedings of the 11th international conference on Intelligent user interfaces, 2006, p. 360362, ISBN 1-59593-287-9.
- [MD06] Lee McCauley and Sidney DMello, ch. MIKI: A Speech Enabled Intelligent Kiosk, pp. 132–144, Springer, 2006, Volume 4133.
- [MdKG08] Louis-Philippe Morency, Iwan de Kok, and Jonathan Gratch, *Context-based recognition during human interactions: automatic feature selection and encoding dictionary*, ICMI '08 Proceedings of the 10th international conference on Multimodal interfaces (Chania, Greece), ACM, 2008, pp. 181–188, ISBN 978-1-60558-198-9.
- [MPWR07] M. F. Macedonio, T. Parsons, B. Wiederhold, and A. A. Rizzo, *Immersion and Physiological Arousal within Panoramic Video-based Virtual Reality*, CyberPsychology and Behavior **10** (2007), no. 4, 508–515, ISSN 1094-9313.
- [PBBR07] T. Parsons, T. Bowerly, J. G. Buckwalter, and A. A. Rizzo, *A controlled clinical comparison of attention performance in children with ADHD in a virtual reality classroom compared to standard neuropsychological methods*, Child Neuropsychology **13** (2007), no. 4, 363–381, ISSN 0929-7049.
- [PE08] M. Powers and P. M. G. Emmelkamp, *Virtual reality exposure therapy for anxiety disorders: A meta-analysis*, Journal of Anxiety Disorders **22** (2008), no. 3, 561–569, ISSN 0887-6185.
- [Pel01] B. Pellom, *Sonic: The University of Colorado continuous speech recognizer*, Tech. report, 2001, Technical Report TR-CSLR-2001-01.

- [PI04] Helmut Prendinger and M. Ishizuka (eds.), *Life-like characters: tools, affective functions, and applications*, Springer, 2004, ISBN 3-540-00867-5.
- [PKN⁺08] T. D. Parsons, P. Kenny, C. A. Ntuen, C. S. Pataki, M. T. Pato, A. A. Rizzo, C. St-George, and J. Sugar, *Objective structured clinical interview training using a virtual human patient*, *Studies in health technology and informatics* **132** (2008), 357–362, ISSN 0926-9630.
- [PR08] T. Parsons and A. A. Rizzo, *Affective Outcomes of Virtual Reality Exposure Therapy for Anxiety and Specific Phobias: A Meta-Analysis*, *Journal of Behavior Therapy & Experimental Psychiatry* **39** (2008), no. 3, 250–261, ISSN 0005-7916.
- [PSB02] D.-P. Pertaub, M. Slater, and C. Baker, *An Experiment on Public Speaking Anxiety in Response to Three Different Types of Virtual Audience*, *Presence* **11** (2002), no. 1, 68–78, ISSN 1054-7460.
- [PSC06] L. Padgett, D. Strickland, and C. Coles, *Case study: Using a virtual reality computer game to teach fire safety skills to children diagnosed with Fetal Alcohol Syndrome (FAS)*, *Journal of Pediatric Psychology* **31** (2006), no. 1, 65–70, ISSN 0146-8693.
- [RBR05] F. D. Rose, B. M. Brooks, and A. A. Rizzo, *Virtual Reality in Brain Damage Rehabilitation: Review*, *CyberPsychology and Behavior* **8** (2005), no. 3, 241–262, ISSN 1094-9313.
- [RCN⁺03] A. Rutten, S. Cobb, H. Neale, S. Kerr, A. Leonard, S. Parsons, and P. Mitchell, *The AS interactive project: single-user and collaborative virtual environments for people with high-functioning autistic spectrum disorders*, *Journal of Visualization and Computer Animation* **14** (2003), no. 5, 233–241, ISSN 1546-427X.
- [RDR08] B. O. Rothbaum, J. Difede, and A. A. Rizzo, *Therapist Treatment Manual for Virtual Reality Exposure Therapy: Post-traumatic Stress Disorder In Iraq Combat Veterans*, Geneva Foundation, 2008.
- [RGG03] A. A. Rizzo, K. Pryor Ghahremani, and S. Gardner, *Human-computer interaction: Theory and practice*, ch. Immersive 360-Degree Panoramic Video Environments, pp. 1233–1237, L.A. Erlbaum: New York, 2003, Vol. 1.
- [RGH⁺01] Jeff Rickel, Jonathan Gratch, Randall Hill, Stacy Marsella, and William Swartout, *Steve Goes to Bosnia: Towards a New Generation of Virtual Humans for Interactive Experiences*, *The Proceedings of the AAAI Spring Symposium on AI and Interactive Entertainment* (Stanford University, CA), 2001.
- [Riv05] G. Riva, *Virtual Reality in Psychotherapy: Review*, *CyberPsychology and Behavior* **8** (2005), no. 3, 220–230, ISSN 1094-9313.
- [RJD⁺07] Andrew B. Raij, Kyle Johnsen, Robert F. Dickerson, Benjamin C. Lok, Marc S. Cohen, Margaret Duerson, Rebecca Rainer Pauly, Amy O. Stevens, Peggy Wagner, and D. Scott Lind, *Comparing Interpersonal Interactions with a Virtual Human to Those with a Real Human*, *IEEE Transactions on Visualization and Computer Graphics* **123** (2007), no. 3, 443–457, ISSN 1077-2626.
- [RK05] A. A. Rizzo and G. Kim, *A SWOT analysis of the field of Virtual Rehabilitation and Therapy*, *Presence: Teleoperators and Virtual Environments* **14** (2005), no. 2, 1–28, ISSN 1054-7460.
- [RKM⁺06] A. A. Rizzo, D. Klimchuk, R. Mitura, T. Bowerly, J. G. Buckwalter, and T. Parsons, *A Virtual Reality Scenario for All Seasons: The Virtual Classroom*, *CNS Spectrums* **11** (2006), no. 1, 35–44, ISSN 1092-8529.
- [RRG⁺09] A. A. Rizzo, G. Reger, G. Grahm, J. Difede, and B. O. Rothbaum, *Post-traumatic stress disorder: Basic science and clinical practice*, ch. Virtual Reality Exposure Therapy for Combat Related PTSD, 2009.

- [RSKM04] A. A. Rizzo, M. T. Schultheis, K. Kerns, and C. Mateer, *Analysis of Assets for Virtual Reality Applications in Neuropsychology*, *Neuropsychological Rehabilitation* **14** (2004), no. 1-2, 207–239, ISSN 0960-2011.
- [RSR02] A. A. Rizzo, Maria T. Schultheis, and Barbara O. Rothbaum, *Ethical issues in clinical neuropsychology*, ch. 12: Ethical issues for the use of virtual reality in the psychological sciences, pp. 243–280, Taylor & Francis, Lisse, NL, 2002.
- [SFW98] D. Stanton, N. Foreman, and P. Wilson, *Virtual reality in clinical psychology and neuroscience*, ch. Uses of virtual reality in clinical training: Developing the spatial skills of children with mobility impairments, pp. 219–232, Amsterdam: IOS Press, 1998.
- [SGJ⁺06] William R. Swartout, Jonathan Gratch, Randall W. Hill Jr., Eduard Hovy, Stacy Marsella, Jeff Rickel, and David Traum, *Toward Virtual Humans*, *AI Magazine* **27** (2006), no. 2, 96–108, ISSN 0738-4602.
- [SGN05] Abhinav Sethy, Panayiotis G. Georgiou, and Shrikanth Narayanan, *Building topic specific language models from webdata using competitive models*, Interspeech-2005 Eurospeech Lisbon, Portugal September 4-8, 2005, pp. 1293–1296.
- [SHJ⁺06] Amy Stevens, Jonathan Hernandez, Kyle Johnsen, Robert Dickerson, Andrew Raji, Cyrus Harrison, Meredith DiPietro, Bryan Allen, Richard Ferdig, Sebastian Foti, Jonathan Jackson, Min Shin, Juan Cendan, Robert Watson, Margaret Duer-son, Benjamin Lok, Marc Cohen, Peggy Wagner, and D. Scott Lind, *The use of virtual patients to teach medical students communication skills*, *The American Journal of Surgery* **191** (2006), no. 6, 806–811, ISSN 0002-9610, Presented at the Association for Surgical Education meeting, March 29, 2005, New York, NY.
- [SPPA⁺04] Susan M. Schneider, Maryjo Prince-Paul, Mary Jo Allen, Paula Silverman, and Deborah Talaba, *Virtual reality as a distraction intervention for women receiving chemotherapy*, *Oncology Nursing Forum* **31** (2004), no. 1, 81–88, ISSN 0190-535X.
- [TFK⁺06] M. Triola, H. Feldmann, A. L. Kaleta, S. Zabar, E. K. Kachur, C. Gillespie, M. Anderson, C. Griesser, and M. Lipkin, *A randomized trial of teaching clinical skills using virtual and live standardized patients*, *Journal of General Internal Medicine* **21** (2006), no. 5, 424–429, ISSN 0884-8734.
- [TMG⁺08] David Traum, Stacy C. Marsella, Jonathan Gratch, Jina Lee, and Arno Hartholt, *Multi-party, Multi-issue, Multi-strategy Negotiation for Multi-modal Virtual Agents*, IVA '08 Proceedings of the 8th international conference on Intelligent Virtual Agents (Tokyo, Japan), Springer, 2008, ISBN 978-3-540-85482-1.
- [TMMK08] Marcus Thiebaut, Stacy Marsella, Andrew N. Marshall, and Marcelo Kallmann, *SmartBody: behavior realization for embodied conversational agents*, *AA-MAS '08 Proceedings of the 7th international joint conference on Autonomous agents and multiagent systems*, 2008, Volume I, pp. 151–158, ISBN 978-0-9817381-0-9.
- [Vir10] Virtually Better, Inc., *Virtually Better Homepage*, www.virtuallybetter.com, 2010, Last visited February 11th, 2011.
- [Wei76] J. Weizenbaum, *Computer Power and Human Reason*, San Francisco: W. H Freeman, 1976.
- [WKFK06] P. L. Weiss, K. Kizony, U. Feintuch, and N. Katz, *Textbook of neural repair and neurorehabilitation*, ch. 13: Virtual Reality in Neurorehabilitation, pp. 182–197, 2006.

- [WOR05] K. Walters, D. Osborn, and P. Raven, *The development, validity and reliability of a multimodality objective structure clinical examination in psychiatry*, *Medical Education* **29** (2005), no. 3, 292298, ISSN 1365-2923.

Citation
Albert Skip Rizzo, Patrick Kenny, and Thomas D. Parsons, <i>Intelligent Virtual Patients for Training Clinical Skills</i> , <i>Journal of Virtual Reality and Broadcasting</i> , 8(2011), no. 3, March 2011, urn:nbn:de:0009-6-29025, ISSN 1860-2037.